

Robust Detection of Ionospheric Irregularities

Todd Walter, Andrew Hansen, Juan Blanch, and Per Enge, *Stanford University*

Tony Mannucci, Xiaoqing Pi, Larry Sparks, and Byron Iijima, *Jet Propulsion Laboratory*

Bakry El-Arini and Roland Lejeune, *MITRE*

Mine Hagen, Eric Altshuler, Rob Fries, and Aleck Chu, *Raytheon*

ABSTRACT

The Wide Area Augmentation System (WAAS) will provide real-time differential GPS corrections and integrity information for aircraft navigation use. The most stringent application of this system will be precision approach, where the system guides the aircraft to within a few hundred feet of the ground. Precision approach operations require the use of differential ionospheric corrections. WAAS must incorporate information from reference stations to create a correction map of the ionosphere. More importantly, this map must contain confidence bounds describing the integrity of the corrections. The confidence bounds must be large enough to describe the error in the correction, but tight enough to allow the operation to proceed. The difficulty in generating these corrections is that the reference station measurements are not co-located with the aviation user measurements. For an undisturbed ionosphere over the Conterminous United States (CONUS), this is not a problem as the ionosphere is nominally well behaved. However, a concern is that irregularities in the ionosphere will decrease the correlation between the ionosphere observed by the reference stations and that seen by the user. Therefore, it is essential to detect when such irregularities may be present and adjust the confidence bounds accordingly.

The approach outlined in this paper conservatively bounds the ionospheric errors even for the worst observed ionospheric conditions to date, using data sets taken from the operational receivers in the WAAS reference station network. As we progress through the current solar cycle and gather more data on the behavior of the ionosphere, many of our pessimistic assumptions will be relaxed. This will result in higher availability while maintaining full integrity.

INTRODUCTION

The nominal or quiet ionosphere above the Conterminous United States (CONUS) is smooth and easily estimated. However, there are times when the ionosphere is more difficult to describe, particularly during geomagnetic and ionospheric storms. Under disturbed conditions smaller scale features may be difficult to observe or estimate. The Wide Area Augmentation System (WAAS) must correct for the users' ionospheric delay errors and place strict confidences on those corrections under all conditions, but its measurements are not co-located with the users'. Therefore, we have to translate the knowledge we gain through our measurements, from their locations, to any possible user location. The mechanism specified in the WAAS Minimum Operational Performance Standards (MOPS) [1] is the vertical ionospheric delay grid. The MOPS specifies transmission of vertical delay values and confidences at discrete grid locations. These confidences must bound the errors not only at the grid locations, but for all interpolated regions between the grid points. In addition we must bound the errors both for the nominal and the disturbed ionosphere.

If we cannot distinguish between nominal and disturbed conditions, then we must always assume disturbed conditions are present. Instead, we would prefer to detect ionospheric irregularities so that we can provide a high level of service during nominal periods. A reduced level of service would only be necessary during periods of detected disturbances. The detection scheme must be extremely robust in order to provide the necessary level of protection. The integrity requirements for precision approach guidance set the probability of hazardously misleading information below 10^{-7} per approach. Therefore, the chance of an undetected ionospheric irregularity must be at a similarly small level.




Figure 1. A two dimensional histogram representing the distribution of differential vertical delays, after a planar fit has been removed, versus IPP separation. This data is for a quiet day. The bar at the right indicates the number of counts per square.

The detection scheme described in this paper verifies the consistency of multiple pierce point samples of the ionosphere. The reference station measurements are fit to a simple linear model. If the model is in error or the input measurements or confidences are incorrect, then statistically the chi-square residual of the fit is affected. If the consistency of the measurements to the model becomes too poor, the local model of the ionosphere cannot be trusted and the confidence bounds must be increased or the grid point may be declared unusable.

In order to perform the fit we must better understand the structure of the nominal ionosphere and then investigate what forms a disturbance might take. The mechanism we use to describe ionospheric behavior is the decorrelation function. This function describes how measurements at one location may differ from the nearby ionosphere.

DECORRELATION FUNCTION

The decorrelation function is used to relate ionospheric measurements made at one location to ionospheric delay values at other locations. This function is determined using post-processed data from reference sites with multiple receivers. Post-processing with the dual-frequency carrier phase removes most of the noise associated with GPS ionospheric code measurements. The redundancy removes errors from the individual receivers. The resulting data has very low noise and no observed receiver artifacts. This so called “supertruth” data is used to probe the structure of the ionosphere. An additive decorrelation function can be formed by differencing adjacent ionospheric measurements and indexing the data by delay difference versus distance. The delays are first




Figure 2. A two dimensional histogram representing the distribution of differential vertical delays, after a planar fit has been removed, versus IPP separation. This data represents a disturbed day.

converted from slant to vertical using the thin shell model assuming a height of 350 km [2].

The data indicate that the expected vertical difference is a linear function of separation distance [3] [4]. These results are consistent with an ionosphere locally well modeled by a plane with some uncertainty about that fit. This so called zeroth order decorrelation was extended to remove a planar fit from the data. After the planar fit is removed, the measurements are differenced again to form the first order decorrelation function. Figures 1 and 2 show two dimensional histograms of differences in vertical delay values, versus Ionospheric Pierce Point (IPP) separation. In this plot a planar fit has first been removed. The shading shows the number of points in each square. Figure 1 has data for a quiet day. Notice that the difference never exceeds 2.2 m throughout.

The observed nominal CONUS first order decorrelation function derived from the data in Figure 1 has form [3] [4]

$$\sigma_{decorr}(el, d_{i,j}) = b \quad (1)$$

where el is the elevation angle of the slant measurement and $d_{i,j}$ is the distance between the two points. Currently the value of b is set to 35 cm. Note that the final form is not dependent on either of the variables. There is no dependency on distance because the planar fit models the nominal CONUS ionosphere extremely well out to distances in excess of 2000 km. It was expected that this decorrelation value would be a function of elevation angle as it contains error in the obliquity factor conversion from slant to vertical [5] [6]. However, the data showed no such dependency. Upon closer investigation it was determined that the dominant error term resulted from low

frequency multipath affecting the post-processed results. Thus, the observed decorrelation function describes the measurement errors rather than the ionosphere. It would be more accurate to say that the decorrelation of the nominal CONUS ionosphere is bounded by 35 cm about a plane.

For disturbed days, the decorrelation function may have the same form but with higher values (b could be more than ten times higher for severe storms). Alternately the decorrelation function could increase with distance, indicating the presence of significant curvature. On these days the nominal decorrelation function (1) will not be an adequate description of the ionosphere. Figure 2 shows data for a disturbed day. Note the change in scale relative to Figure 1.

ESTIMATION OF DELAY AND ERROR

We wish to determine the delay and confidences for an Ionospheric Grid Point (IGP) given N Ionospheric Pierce Point (IPP) measurements and confidences within a certain cutoff radius around that IGP. We will use vector notation to denote the delay values, $\mathbf{I}_{v,IPP}$, and the measurement variances, $\sigma_{I,IPP}^2$, as

$$\mathbf{I}_{v,IPP} = \begin{bmatrix} I_{v,IPP_1} \\ I_{v,IPP_2} \\ \vdots \\ I_{v,IPP_N} \end{bmatrix}, \quad \sigma_{I,IPP}^2 = \begin{bmatrix} \sigma_{I,IPP_1}^2 \\ \sigma_{I,IPP_2}^2 \\ \vdots \\ \sigma_{I,IPP_N}^2 \end{bmatrix} \quad (2)$$

where the element σ_{I,IPP_i}^2 is the bounding variance describing noise, multipath, and bias uncertainty for the i^{th} IPP. The full weighting matrix for these IPPs is

$$\mathbf{W}^{-1} = \begin{bmatrix} \sigma_{I,IPP_1}^2 + \sigma_{decorr}^2 & \sigma_{bias,1,2} & \cdots & \sigma_{bias,1,N} \\ \sigma_{bias,1,2} & \sigma_{I,IPP_2}^2 + \sigma_{decorr}^2 & \cdots & \sigma_{bias,2,N} \\ \vdots & \vdots & \ddots & \vdots \\ \sigma_{bias,1,N} & \sigma_{bias,2,N} & \cdots & \sigma_{I,IPP_N}^2 + \sigma_{decorr}^2 \end{bmatrix} \quad (3)$$

where

$$\sigma_{bias,i,j} \equiv \begin{cases} 0 & i \text{ and } j \text{ share no common satellites or receivers} \\ \frac{\sigma_{bias,sat}^2}{F(el_i)F(el_j)} & i \text{ and } j \text{ share a common satellite} \\ \frac{\sigma_{bias,receiver}^2}{F(el_i)F(el_j)} & i \text{ and } j \text{ share a common receiver} \end{cases} \quad (4)$$

as individual IPP measurements are assumed to be otherwise uncorrelated.

The vertical ionospheric delay estimate at the IGP will be a linear combination of the IPP measurements. Provided we have sufficient measurements, this can be done to arbitrary order. The decorrelation data, however, indicates that a first order (planar) fit is most appropriate for quiet days. A zeroth order (weighted average) fit fails to model observed ionospheric structure and no second order (quadratic) terms are evident in the data for nominal days. In addition, over the radii of interest, it would be very difficult to distinguish quadratic terms from planar ones.

We will use a local cartesian frame whose origin is at the IGP and whose x -axis aligns with the East direction at the IGP and whose y -axis aligns with the North. The states are the vertical delay at the IGP, the slope in the East direction and the slope in the North direction. In the region about the IGP the ionosphere is estimated by

$$\hat{I}_{v,IGP}(x,y) = \hat{a}_0 + \hat{a}_1 \cdot x + \hat{a}_2 \cdot y \quad (5)$$

The observation matrix is given by

$$\mathbf{G} = \begin{bmatrix} 1 & \mathbf{d}_{IPP_1,IGP} \cdot \hat{\mathbf{E}} & \mathbf{d}_{IPP_1,IGP} \cdot \hat{\mathbf{N}} \\ 1 & \mathbf{d}_{IPP_2,IGP} \cdot \hat{\mathbf{E}} & \mathbf{d}_{IPP_2,IGP} \cdot \hat{\mathbf{N}} \\ \vdots & \vdots & \vdots \\ 1 & \mathbf{d}_{IPP_N,IGP} \cdot \hat{\mathbf{E}} & \mathbf{d}_{IPP_N,IGP} \cdot \hat{\mathbf{N}} \end{bmatrix} \quad (6)$$

We can solve for the planar coefficients

$$[\hat{a}_0 \quad \hat{a}_1 \quad \hat{a}_2] = \left[(\mathbf{G} \cdot \mathbf{W} \cdot \mathbf{G}^T)^{-1} \cdot \mathbf{G} \cdot \mathbf{W} \cdot \mathbf{I}_{v,IPP} \right]^T \quad (7)$$

and find the delay estimate at the IGP given by

$$\hat{I}_{v,IGP} = \hat{a}_0 = [1 \quad 0 \quad 0] \cdot \left[(\mathbf{G} \cdot \mathbf{W} \cdot \mathbf{G}^T)^{-1} \cdot \mathbf{G} \cdot \mathbf{W} \cdot \mathbf{I}_{v,IPP} \right] \quad (8)$$

and the formal error on the estimate is

$$\hat{\sigma}_{\hat{I}_{v,IGP}}^2 = \left[(\mathbf{G} \cdot \mathbf{W} \cdot \mathbf{G}^T)^{-1} \right]_{1,1} = \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}^T \cdot (\mathbf{G} \cdot \mathbf{W} \cdot \mathbf{G}^T)^{-1} \cdot \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} \quad (9)$$

Thus, given a set of measurements we can generate the vertical correction and confidence factor for a vertical delay at that grid point under nominal conditions. Unfortunately, the ionosphere is not always in its

nominal state and we must protect all slant delays which may be interpolated from this IGP. The following sections address these additional terms, which must be incorporated in our final confidence.

IRREGULARITY DETECTION

We have based our delay and confidence estimation on a local planar model with uncertainties bounded by a 35 cm standard deviation. Therefore, we will define irregular ionospheric behavior to be any condition that cannot be accurately described by such a model. This may not follow conventional descriptions of ionospheric events, which seek to distinguish events by their physical cause (storms, irregularities, scintillation, etc.). Here, we are only interested in events which invalidate the confidence calculation (9). A storm which causes a steep gradient, but is well modeled and bounded by (8) and (9) requires no special attention from WAAS. However, an event which violates our confidence modeling, must be recognized and treated. We generically describe all such events as irregularities and must formulate a means to detect them.

Fortunately, there is a well established mechanism for determining if measurements are consistent with the model. This test is known as a chi-square consistency check or “goodness-of-fit” test. Since we know the model, the measurement values, and their variances, we can use the expected distribution of the residual error to determine if they are consistent.

We investigated other irregularity detectors based on spatial gradients. It has been observed that poor grid modeling correlates strongly with large gradients. The gradient detectors performed equivalently (and in some cases better) on the data examined, but in the end, the chi-square had the best analytical connection between the threshold and the resulting error, and was chosen for that reason.

The planar model can be evaluated around the IGP to provide estimates for each IPP

$$\hat{\mathbf{I}}_{v,IPP} \equiv \begin{bmatrix} 1 \\ \mathbf{d}_{IPP,IGP} \cdot \hat{\mathbf{E}} \\ \mathbf{d}_{IPP,IGP} \cdot \hat{\mathbf{N}} \end{bmatrix}^T \cdot \left[(\mathbf{G} \cdot \mathbf{W} \cdot \mathbf{G}^T)^{-1} \cdot \mathbf{G} \cdot \mathbf{W} \cdot \mathbf{I}_{v,IPP} \right] \quad (10)$$

The chi-square statistic is defined as

$$\begin{aligned} \chi^2 &= \mathbf{I}_{v,IPP}^T \cdot \mathbf{W} \cdot \left\{ \mathbf{I} - \mathbf{G}^T \cdot (\mathbf{G} \cdot \mathbf{W} \cdot \mathbf{G}^T)^{-1} \cdot \mathbf{G} \cdot \mathbf{W} \right\} \cdot \mathbf{I}_{v,IPP} \\ &= \mathbf{I}_{v,IPP}^T \cdot \mathbf{W} \cdot (\mathbf{I} - \mathbf{P}) \cdot \mathbf{I}_{v,IPP} \end{aligned} \quad (11)$$

The chi-square statistic is a reliable indicator of the “goodness of fit.” If either the input variances or the model are incorrect, this measure will be affected statistically. A parameter in the expected distribution is the number of degrees of freedom of the statistic. This is simply the number of measurements minus the degrees of freedom in the model (three for 1st order). Given this parameter and an allowable false alarm rate, P_{fa} , we can calculate a threshold value for the chi-square using its known distribution [7] (see third column of Table 1). If the chi-square statistic is above this value, it is likely that there is a problem with the model, the measurements, or both. The difficulty is detectability. A chi-square value below the threshold does not automatically guarantee that the model is valid.

ν	$\chi_{0.001}^2$	$\chi_{0.999}^2$	R_{irreg}^2
2	0.2102	20.5150	97.5917
3	0.3811	22.4577	58.9339
4	0.5985	24.3219	40.6385
5	0.8571	26.1245	30.4799
6	1.1519	27.8772	24.2000
7	1.4787	29.5883	20.0091
8	1.8339	31.2641	17.0483
9	2.2142	32.9095	14.8629
10	2.6172	34.5282	13.1927
11	3.0407	36.1233	11.8800
12	3.4827	37.6973	10.8242
13	3.9416	39.2524	9.9584
14	4.4161	40.7902	9.2367
15	4.9048	42.3124	8.6266
16	5.4068	43.8202	8.1046
17	5.9210	45.3147	7.6532
18	6.4467	46.7970	7.2591
19	6.9830	48.2679	6.9122
20	7.5292	49.7282	6.6047
21	8.0849	51.1786	6.3302
22	8.6493	52.6197	6.0837
23	9.2221	54.0520	5.8611
24	9.8028	55.4760	5.6592
25	10.3909	56.8923	5.4752
26	10.9861	58.3012	5.3068
27	11.5880	59.7031	5.1522

Table 1. For different degrees of freedom, ν ; the lower bound, $\chi_{0.001}^2$ ($P_{md} = 10^{-3}$); threshold, $\chi_{0.999}^2$ ($P_{fa} = 10^{-3}$); and inflation factor, R_{irreg}^2 , are given. Since a planar fit is performed, $\nu = N-3$, where N is the number of IPPs.

The chi-square test does protect against certain error conditions, such as large errors on any single measurement or moderate errors on all of the measurements. However, if the errors are only slightly larger than expected, the confidence bound predicted by (9) will be optimistic but there is some probability that the chi-square statistic will not exceed the threshold. When the actual variances of the errors are larger than expected (perhaps because the assumed decorrelation function does not adequately describe the current state of the ionosphere), we must determine the probability of missing this event and inflate our bound as necessary.

WORST UNDETECTED DISTRIBUTION

There are many ways to model a failure. Non-planar behavior in the ionosphere will result in unmodeled biases in the residuals. Alternately, if some of the true variances increase beyond their assumed values, the chi-square distribution will increase, but still have some likelihood of being below the threshold. A more generic failure mode could contain a combination of deterministic biases and increased variances distributed in various ways among the IPPs. This approach would require use of the non-central chi-square distribution. [7] We will restrict ourselves to a pessimistic, but tractable model in which all variances increase uniformly by the same factor, R_{irreg}^2 .

We can use the known distribution of this failure model to determine the missed detection probability. This probability, P_{md} , is given by the integration of the new distribution's Probability Density Function (PDF) below the threshold. Alternately, we can determine the required value of R_{irreg}^2 to achieve the specified P_{md} . From (11) we can see that a constant term multiplying each variance results in a new chi-square statistic which is equal to the original value divided by this factor. The resulting distribution is essentially unchanged, except that it is stretched along the chi-square axis (see Figure 3).

To determine the numerical value of this factor, we must concentrate on the lower bounds of the distribution. We must have less than P_{md} of the new distribution below the threshold. The factor, R_{irreg}^2 , can be determined by taking the ratio of the upper threshold to the lower bound

$$R_{irreg}^2(P_{fa}, P_{md}) = \frac{\chi_{1-P_{fa}}^2}{\chi_{P_{md}}^2} \quad (12)$$

where χ_p^2 denotes the inverse of the Cumulative Density Function (CDF), or the chi-square value that has p of the distribution below it.




Figure 3. The blue trace shows the expected distribution for the chi-square variable. The black line represents the threshold listed in Table 1. The red trace corresponds to the worst case undetectable distribution when all the variances are multiplied by 6.69 with a corresponding $P_{md} = 10^{-1}$.

Figure 3 depicts an example using nine degrees of freedom with P_{fa} set equal to 10^{-3} and P_{md} equal to 10^{-1} . The expected PDF is shown in blue. In this example, $\chi_{0.999}^2 = 27.9$, and $\chi_1^2 = 4.17$, therefore $R_{irreg}^2 = 6.69$. The chi-square distribution with the variances inflated by R_{irreg}^2 is shown in red. As expected, only 10% of this new distribution lies below the threshold value. Thus, under our failure model, this is the worst distribution that could escape detection with probability P_{md} . We call this the worst undetected distribution. In this example, the chi-square check verifies that the actual distributions do not uniformly have variances 6.69 times larger than the expected values with a $P_{md} = 10^{-1}$.

Because this is the worst distribution the chi-square check protects against, we must generate our confidence bounds according to it. The variance in (9) must be multiplied by R_{irreg}^2 to ensure that it guards against all distributions that pass the chi-square check. Thus, we form the estimate and perform the chi-square check assuming the quiet decorrelation function. Then, if the fit passes this check, we recalculate the confidence using the scaled distribution. The scaled distribution is the best that we can protect to the desired level of certainty.

The bounding variance for a raypath through that IGP is given by

$$\sigma_{\text{bound, IGP}}^2 = R_{irreg}^2 \cdot \left[(\mathbf{G} \cdot \mathbf{W} \cdot \mathbf{G}^T)^{-1} \right]_{1,1} + R_{irreg}^2 \cdot \sigma_{\text{decorr}}^2 \quad (13)$$

where the first term accounts for the uncertainty in estimating the plane, and the second term describes the

inherent uncertainty about that plane. Translation of the confidence from the grid point to any point in the cell is addressed later in the paper.

VALIDATION RESULTS

We also used the supertruth data to validate the algorithms described above. We centered a planar fit around every IPP in the supertruth sets. The IPP at the center was excluded from the fit. In this way we had an independent truth reference and could use the other data points to verify how well they predicted the vertical delay at the excluded IPP. The algorithm was executed creating a predicted delay (8), a bounding confidence (9), and a chi square indication of the “goodness-of-fit” (11).

Figure 4 shows the distribution of errors divided by the square root of the bounding variance, here given by

$$\frac{I_{v,IPP} - \hat{I}_{v,IPP}}{\sqrt{\left[(\mathbf{G} \cdot \mathbf{W} \cdot \mathbf{G}^T)^{-1} \right]_{1,1}} + \sigma_{decorr}^2} \quad (14)$$

No R_{irreg}^2 term is used yet as we know that this is a quiet ionospheric day and it is well modeled by the nominal decorrelation function.

If the predicted vertical delays and bounding variances were accurate, the resulting distribution should follow a zero-mean unit-variance gaussian which is shown for reference by the red line. As can be seen, the confidence estimate is conservative; it dramatically overbounds at the tails. At no time during the whole day is the error greater than 2.5 times the predicted standard deviation. The data is very clean and gaussian in nature, but with clipped tails. Although the algorithm was designed to have a false alarm rate of 10^{-3} , the chi-square value never exceeded its threshold value during this day.

If the ionosphere always behaved in this manner, the irregularity detector and the R_{irreg}^2 term would be unnecessary. Unfortunately, there are times when the ionosphere is disturbed. During such events, the planar approximation or decorrelation description may break down. Figure 5 shows the same type of plot but for data from April 6 and 7, and July 15 and 16 of 2000. These days represent the worst CONUS ionospheric data for the current solar peak. Now we see that the nominal confidence bounds are inadequate. Errors can exceed 50 times the predicted standard deviation. It is for this reason that we have implemented the chi-square check and the variance inflation.




Figure 4. The histogram of prediction error divided by nominal predicted error (14) for December 6th, 1999, a nominal day.




Figure 5. The histogram of prediction error divided by the nominal predicted error (14) for April 6&7 and July 15&16, 2000, four severely disturbed days.

Figure 6 shows the data now for only those points whose chi-square statistic was below the threshold. Data that failed the check are now excluded. As expected, this improves the distribution but does not remove all outliers. The true error can still be greater than 10 times the predicted standard deviation. Data with worse than nominal distribution still has some chance to pass the chi-square test. If we apply the R_{irreg}^2 inflation term, the new distribution becomes

$$\frac{I_{v,IPP} - \hat{I}_{v,IPP}}{R_{irreg} \cdot \sqrt{\left[(\mathbf{G} \cdot \mathbf{W} \cdot \mathbf{G}^T)^{-1} \right]_{1,1}} + \sigma_{decorr}^2} \quad (15)$$

as shown in Figure 7. This final set is well described by the inflated variance. This full process of calculating the nominal estimates, applying a chi-square check and




Figure 6. The histogram of the error ratio (14) for the disturbed days when data that failed the chi-square check is excluded.




Figure 7. The histogram of the error ratio (15) for the disturbed days when data that failed the chi-square check is excluded and the predicted variance is inflated by R_{irreg}^2 .

inflating the bounding confidence for those that pass, is necessary to protect the user.

Figure 7 contains some outliers which do not match the distribution of the majority of the points. These points extend out to four standard deviations, still safe, but uncomfortably large. Upon closer examination it was determined that the error was not very large (< 5 m), but that the confidence had grown too small. For these test runs the off-diagonal elements were left out of the weighting matrix. Operationally these bias correlation terms will be included, but here their omission allowed the formal error to become too small. Including these terms will worsen availability, but result in a safer distribution.

Values that fail the chi square test must be treated separately. Failed IGP's could be set to "Do Not Use" or




Figure 8 The location of all IPPs within 10° of the IGP are shown by the blue asterisks. The mesh is the evaluation of (16) for locations within the 5° square surrounding the IGP. Also shown is a projection of this mesh at the bottom. The maximum uncertainty of 0.35 m within this square occurs at the Southwest corner. The σ_{GIVE} would need to be at least this large.

could be recalculated using a different model. Another possibility is to exploit the known physical limitations of the ionosphere. The largest vertical delay observed in CONUS is well below 64 m. Therefore, a one sigma bound of ~ 12 m would protect any true error that fits within the MOPS message format to within 5.33 sigma or to 10^{-7} . Thus, provided separate arguments or monitoring can be made to protect against correlated measurements on multiple satellites, it could be argued that the largest GIVE index in the MOPS would always be safe.

SPATIAL ERROR

The broadcast variance, termed σ_{GIVE}^2 , must bound the worst error that a user may experience in the four surrounding cells after interpolating from the corresponding IGP's. For most of the grid, this problem can be simplified by instead concentrating on the 5° by 5° square centered on the IGP. The value of σ_{GIVE}^2 must be at least as large as the largest value in this square. For the planar model, the formal error will have a global minimum near the weighted centroid of the measurements and will increase as one moves away from this point. Therefore, the maximum uncertainty over the central square will be at the farthest distance from the centroid. This will correspond to one of the four corners as illustrated in Figure 8. Thus the GIVE value can be determined from the corner with greatest uncertainty.

The formal error for the IPP estimate is

$$\hat{\sigma}_{\hat{I}_{v,IPP}}^2 \equiv \begin{bmatrix} 1 \\ \mathbf{d}_{IPP,IGP} \cdot \hat{\mathbf{E}} \\ \mathbf{d}_{IPP,IGP} \cdot \hat{\mathbf{N}} \end{bmatrix}^T \cdot (\mathbf{G} \cdot \mathbf{W} \cdot \mathbf{G}^T)^{-1} \cdot \begin{bmatrix} 1 \\ \mathbf{d}_{IPP,IGP} \cdot \hat{\mathbf{E}} \\ \mathbf{d}_{IPP,IGP} \cdot \hat{\mathbf{N}} \end{bmatrix} \quad (16)$$

$$\sigma_{GIVE}^2 = \left(\max \left\{ \begin{array}{l} \hat{\sigma}_{\hat{I}_{v,IGP \text{ Lat}+2.5^\circ \text{ Lon}+2.5^\circ}}^2, \hat{\sigma}_{\hat{I}_{v,IGP \text{ Lat}+2.5^\circ \text{ Lon}-2.5^\circ}}^2, \\ \hat{\sigma}_{\hat{I}_{v,IGP \text{ Lat}-2.5^\circ \text{ Lon}+2.5^\circ}}^2, \hat{\sigma}_{\hat{I}_{v,IGP \text{ Lat}-2.5^\circ \text{ Lon}-2.5^\circ}}^2 \end{array} \right\} \right. \\ \left. + \sigma_{decorr}^2 \right) \cdot R_{irreg}^2 \quad (17)$$

The algorithms described above are the complete set of steps from pierce point measurement to GIVE calculation.

The availability is a strong function of cutoff radius. The center of the country is well served by smaller radii (~500 km), but the edges will likely require greater radii (1000 - 2000 km) to support sufficient measurements. Additionally the uncertainty at the edges will be high as most of the measurements will be to one side of the IGP resulting in poorer observation geometry. When combined with the lower overall number of measurements, the edge IGPs will have particularly poor coverage. Larger radii offer the advantage of a larger number of IPPs which lower the formal error and enhance the consistency check. The disadvantages of larger radii are increased computational load and greater averaging of the ionosphere. In addition, small irregularities are less likely to trip the detector.

UNDERSAMPLED THREAT

The detector outlined above works very well provided the irregularity is sampled by several IPPs. This is expected to be true as it has been observed that irregularities occur on multiple scales simultaneously. If the ionosphere is disturbed in one region, it is likely disturbed around that region as well. Thus, even if there is a significant but very localized disturbance, it is expected the surrounding ionosphere will exhibit some irregular behavior. For this reason, even though the reference station IPPs may not sample the worst location in the ionosphere, if the IPPs are sufficiently dense, the surrounding measurements will trigger the detector. If the IPPs are not particularly dense, as at the edge of coverage, additional protection may be needed.

We are in the process of defining a threat model which describes the worst case ionosphere as a function of scale length, given that the chi-square detector has not tripped. This investigates disturbances that occur in isolation, that is, without the accompanying structure that guarantees tripping of the detector. Such disturbances will be much




Figure 9. A two dimensional histogram representing the distribution of differential vertical delays, after a planar fit has been removed, versus IPP separation. This data is for a severe storm day where the chi-square detector has been used to remove locally disturbed ionosphere. Compare to Figures 1 and 2.

smaller in magnitude than features that can develop during fully disturbed conditions. The magnitude is expected to scale with the size of the region in which they are confined.

Undersampled threats would escape any detector and are not limited to the chi-square check. If we have not sampled a threat, we cannot identify it. Threats of this nature must be treated by off-line analysis, where they are rigidly bounded and then incorporated as though always present. The extent of these threats is determined by the density of measurements provided by the reference network. We must ensure that our network is sufficiently dense as to keep the overall effect small.

As an example of the processing required, we used the validation processing previously described, to exclude certain points from a new calculation of the decorrelation function. Here, if the surrounding ionosphere, excluding the specific IPP, tripped the chi-square detector, then the specific IPP is not used to form the new decorrelation plot. Figure 9 shows the decorrelation histogram for the July 15, 2000 severe storm using only IPPs whose surrounding ionosphere did not trip the detector. As can be seen this decorrelation plot closely resembles the quiet data in Figure 1. Thus, none of the IPPs for the worst day sampled an isolated event inconsistent with the nominal assumptions. We will continue to process more storm days as they occur, and refine the quantitative model as a function of scale length. We are also investigating other metrics to ensure that the threat is properly characterized.

RATE OF CHANGE

The ionospheric corrections and confidences are expected to be sent at the slow update rate of once every five minutes. Therefore we need to ensure that the bounding variance is applicable over the lifetime of the message. In addition we need to continue to monitor already broadcast information to detect sudden changes in the ionosphere. This is described in another component of the threat model that examines the time rate of change of the ionosphere. Knowing how quickly the ionosphere can change allows us to determine how much the variance needs to be increased as well as the update rates for testing the validity of old messages.

As in the previous section we want to know the worst rate of change just prior to tripping the chi-square detector. We know that the highest rates of change occur during severely disturbed periods. However, if the chi-square detector has already tripped, it is providing protection against these maximum rates of change. We need to determine the maximum rate of change when the chi-square has not tripped. This latter value should be much smaller than if we did not exploit the chi-square check. When fully detailed, the threat model will provide histograms of the change in vertical delay values over given periods of time, analogous to the spatial decorrelation data already presented. From this information we will be able to set the detector test rates and the additional variance term which will could be added to Equation (17).

PROTOTYPE RESULTS

The algorithm described in this paper has been implemented on the prototype of the operational system. This prototype uses the actual WAAS reference station network and real-time estimates of the noise levels. Figure 10 shows the results for July 11, 2000, a quiet day. As can be seen, the true distribution is quite similar in appearance to the theoretical expectation (red line). The most noticeable discrepancy results from the overbounding of the input noise variances (2). The conservative real-time estimates of these values results in a chi-square value which is smaller than would be obtained if the variances were exactly known, thus shifting the distribution to the left of theory. Therefore, this figure indicates that the computed formal error is also conservative, as the input estimated variances are clearly larger than the true values. The shape of the distribution follows expectation and the false alarm rate for this quiet day is better than required. While the integrity analysis needs to be repeated for the prototype code, the combination of off-line safety




Figure 10. Histogram of the chi-square statistic for a quiet day on the prototype system. The red line indicates the theoretical result. Since the noise estimates tend to be conservative, the observed distribution is shifted to the left of the theory.

validation of the “supertruth” data and the low false alarm rate of the real-time prototype code give us a high degree of confidence for the final stages of validation.

CONCLUSIONS

It has been shown that the nominal ionosphere over CONUS is well behaved and easily modeled by the MOPS grid. A high level of performance is achievable under quiet conditions. However, we have also observed conditions which are not well modeled by the grid. A definition of irregular behavior has been provided. This definition is aimed at the WAAS MOPS broadcast mechanism and the internal ionospheric model. A method for identifying irregular behavior has been derived and validated. This method conservatively assumes that the ionosphere is always distributed in the worst undetectable manner unless the chi-square statistic exceeds a threshold. When the chi-square value is larger than the threshold, the bounding variance is increased even further. Thus integrity is firmly maintained.

We have shown that for quiet days this method is exceedingly conservative. However, on severely disturbed days the worst case assumptions are approached and therefore necessary to protect the user. This algorithm preserved integrity even on the worst observed storm days of the current solar cycle. This algorithm was analytically derived before the storms took place and performed as expected for these severe disturbances. The single empirical parameter, σ_{decorr} , was determined using data from quiet days. Finally, we have implemented this algorithm on the prototype of the operational system and found it to match expectations. In particular, the false

alarm rates are better than the requirement for quiet days. Combined with the successful off-line integrity analysis, we feel that this algorithm is capable of meeting the stringent requirements for providing differential corrections and bounds for vertical aircraft guidance.

ACKNOWLEDGMENTS

This work was sponsored by the FAA GPS Product Team (AND-730). We would also like to thank our colleagues participating in the WAAS Integrity Performance Panel (WIPP) for their guidance and numerous contributions.

REFERENCES

- [1] RTCA Special Committee 159 Working Group 2, "Minimum Operational Performance Standards for Global Positioning System / Wide Area Augmentation System Airborne Equipment," RTCA Document Number DO-229B, October 1999.
- [2] Global Positioning System Standard Positioning Service Signal Specification, June 1995.
- [3] Hansen, A., Peterson, E., Walter, T., and Enge, P., "Correlation Structure of Ionospheric Estimation and Correction for WAAS," in proceedings of ION NTM, Anaheim, CA, January 2000.
- [4] Hansen, A., Walter, T., Blanch, J., and Enge, P., "Ionospheric Spatial and Temporal Correlation Analysis for WAAS: Quiet and Stormy," in proceedings of ION GPS, Salt Lake City, UT, September 2000.
- [5] Sparks, L., Iijima, B.A., Mannucci, A.J., Pi, X., and Wilson, B.D., "A New Model for Retrieving Slant TEC Corrections for Wide Area Differential GPS," in proceedings of ION NTM, Anaheim, CA, January 2000.
- [6] Conker, R. and El-Arini, B., "A Novel Approach for an Ionospheric Obliquity Process Responsive to Azimuthal Variation," in proceeding of ION Annual Meeting, Denver, CO, June 1998.
- [7] Abramowitz, M. and Stegun, I. E., *Handbook of Mathematical Functions*, 1972.
- [8] Pi, X., Mannucci, A. J., Lindqwister, U. J., and Ho, C. M., "Monitoring of Global Ionospheric Irregularities Using the Worldwide GPS Network," *Geophys. Res. Lett.*, **24**, 2283, 1997.